

(19)



Europäisches Patentamt

European Patent Office

Office européen des brevets



(11) Publication number:

0 477 448 A1

(12)

EUROPEAN PATENT APPLICATION

(21) Application number: 90310699.5

(51) Int. Cl.5: H04L 12/26

(22) Date of filing: 28.09.90

(43) Date of publication of application:
01.04.92 Bulletin 92/14(84) Designated Contracting States:
DE FR GB IT(71) Applicant: Hewlett-Packard Company
Mail Stop 20 B-O, 3000 Hanover Street
Palo Alto, California 94304(US)(72) Inventor: Phaal, Peter
25 The Valls
Bradley Stoke, Bristol(GB)(74) Representative: Squibbs, Robert Francis et al
Hepworth, Lawrence, Bryer & Bizley, Lewins
House, Lewins Mead
Bristol BS1 2NN(GB)

(54) Network monitoring device and system.

(57) A network monitoring device (12) is provided for monitoring the activity on a network carrying message packets each of which contains source and destination addresses. The monitoring device (12) comprises a network interface for sending and receiving message packets carried on the network, and processing means operative to collect and process data about the packets received by the network interface. In order to minimise processor memory requirements for the monitoring device (12), only randomly selected ones of the packets detected by the network interface are processed by the processing means of the device. Preferably, the monitoring

device (12) is further simplified by arranging for the data collected on the randomly sampled packets to be transmitted to a central measurement station (13) for analysis. As a result, the only processing required to be done by the monitoring device (12) is the construction of collected-data packets for transmission. A network monitoring system can advantageously be provided by using a number of such monitoring devices (12) each associated with a respective logical segment of the network and each forwarding collected-data packets to a central measurement station (13).

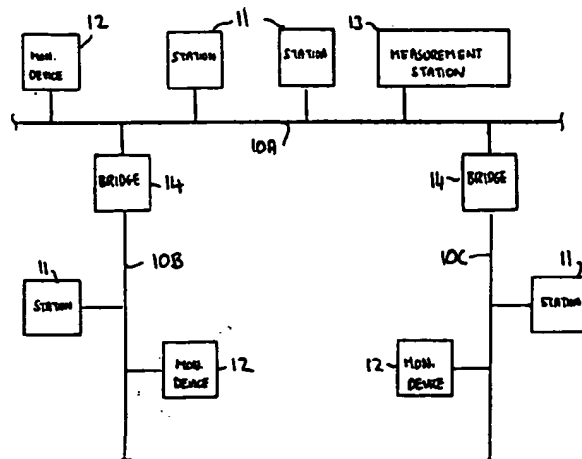


FIG. 1

EP 0 477 448 A1

The present invention relates to a network monitoring device for monitoring the activity on a network carrying message packets each of which contains source and destination addresses; the invention further relates to a network monitoring system utilizing such devices.

Currently most of the information used to manage networks comes from monitoring network devices such as bridges and routers whose primary function is to control the passage of packets between sections of the network. These devices generally provide information about their configuration and some interface statistics. The interface statistics are usually in the form of counts of different types of packet processed by the devices. These counts include the total number of correctly transmitted and received packets and the total number of errored packets, possibly broken down into categories (such as CRC, collision, runt, jabber etc).

The problem with these counts is that although they may be used to indicate a problem (such as an excessive packet collision rate) they do little to isolate the cause.

At this point a traffic matrix is often useful. It breaks down the packet count into the contributions of each station on the network. For example, if the total packet count was high it would be interesting to see which pairs of network stations were communicating and their relative contributions to the packet count. It is only with this information that a decision can be made as to whether to move a station, add capacity or duplicate a service.

Calculating traffic matrices is an expensive operation that involves decoding every packet on the network. In addition large amounts of memory are used to build up the table. It is not surprising then that network devices do not usually provide traffic matrices, it would be too expensive and it would impair their primary function.

Currently if one wants to build up a traffic matrix an instrument is used. Such instruments generally comprise receive means for detecting and receiving message packets carried on the network, and processing means operative to collect and process data about packets received by the receive means, this data including the source and destination addresses contained in the packet. Instruments are usually too costly to leave in place and so someone would have to connect an instrument to the required place, collect the traffic matrix and then repeat the exercise for the next location to be monitored. This can be a difficult and time consuming task in a widely distributed network.

A further problem is that as networks become faster it is becoming harder to design instruments that will keep up with the data rate on the network.

According to the present invention, there is

provided a network monitoring device for monitoring the activity on a network carrying message packets each of which contains source and destination addresses, the monitoring device comprising receive means for detecting and receiving message packets carried on the network, and processing means operative to collect and process data about packets received by the receive means, this data including the source and destination addresses contained in the packet, characterized in that

- the monitoring device further comprises sampling means for selecting some only of the packets detected by the network interface, and
- the processing means is operative to process data collected in respect of the selected packets only.

Using sampling techniques allows traffic matrices to be constructed without the processor and/or memory requirements of a conventionally produced traffic matrix. This would allow traffic matrices to be collected by devices like bridges and routers acting as monitoring devices and also allows instruments to be constructed that are cheap enough to leave in place. Of course, using only selected packets, rather than all packets, to generate a traffic matrix yields an approximate traffic matrix, but over the period of an hour, say, the estimated traffic matrix can be a very good approximation of the actual traffic matrix.

The sampling means may carry out its selection of packets in a deterministic manner either on the basis of selecting every nth packet (for example, every hundredth packet) or on the basis of selecting the first packet detected after a fixed interval from the last selected packet. However, such a selection process will only enable a realistic traffic matrix to be constructed from the selected packets if the traffic itself is random in nature; where this is not the case, such a deterministic selection process could lead to significant distortions of the traffic matrix. For example, in an extreme case a network station might be programmed to transmit a packet with the same frequency as sampling is carried out by the network monitor; in this event, the packets transmitted by the network station would either be totally missed by the sampling process or would be the only packets selected by the sampling process.

Accordingly, the sampling means preferably effects its selection of packets in a statistically random manner. Such a random selection can be carried out on the basis of elapsed time since the previous packet selection or on the number of packets detected by the network interface.

Preferably, the monitoring device further comprises transmit means for transmitting packets over a network, and the processing means is operative

to process the data collected on said selected packets by forming collected-data packets in each of which is included the data collected in respect of one or more of said selected packets, the processing means being operative to cause said transmit means to transmit each said collected-data packet over the network for remote receipt and further processing. Using such a monitoring device it becomes possible to construct a low cost monitoring system by distributing the low cost monitoring devices about the network to be monitored and arranging for them to transmit their data samples, in the form of collected-data packets, back to a single measurement station for processing (for example, to construct traffic matrices). This arrangement minimizes the processor and memory requirements for each monitoring device and concentrates these expensive resources in a central station.

Generally, the network to be monitored will include one or more bridges and/or routers separating the network into a plurality of sections. In this case, each section preferably has its own associated sampling monitoring device; the measurement station can then not only produce traffic matrices for each individual network section but can also determine the topology of the network in terms of network sections to which monitoring devices are connected.

The transmission of the collected data packets to the measurement station can be effected either over the network being monitored or over a separate network to which the transmit means of the monitoring devices are connected.

The sampling monitoring device can take the form of a stand-alone item, a card in a networked computer (using the computer for power and a slot only), a modified bridge or router, or a process running on a processor of a connected network station.

A network monitoring system utilizing a number of sampling monitoring devices embodying the invention, will now be described by way of non-limiting example with reference to the accompanying diagrammatic drawings, in which:

Figure 1 is an overall diagram of a network to which a measurement station and a number of sampling monitoring devices have been connected to form a network monitoring system;

Figure 2 is a diagram illustrating the general form of a data packet transmitted over the Figure 1 network;

Figure 3 is a block diagram of a sampling monitoring device of Figure 1; and

Figure 4 is a flow chart illustrating the main interrupt service routine run by a controlling microprocessor of the

Figure 3 device.

Figure 1 illustrates a typical local area network in which a plurality of stations 11, 12, and 13 are interconnected via cable segments 10A, 10B, and 10C. The network is divided into three logical segments by bridges 14 that connect respective ones of the cable segments 10B, 10C to the cable segment 10A. As is well known in the art, the bridges serve to filter traffic passing between the network segments, such that messages originating from a particular segment and destined for a station on the same segment are not passed through the bridge or bridges 14 to the other segments whereas messages originating in one segment and intended for another one are allowed across the bridge.

In the illustrated local area network, messages between the stations 11, 12 and 13 are transmitted in the form of packets that are broadcast over the network. Typically a packet will have the form illustrated in Figure 2 with a packet header 15 containing a source address (the address of the station sending the packet) and a destination address (the address of the station intended to receive the packet), and an information field 16 containing the data to be passed to the receiving station and normally including error checking codes. Depending on the particular packet format being used, other fields may also be present; thus, for example, there may be a CRC (cycle redundancy check) field covering both the packet header and information field.

The Figure 1 network may, for example, be an Ethernet network well known to persons skilled in the art.

The network of Figure 1 is arranged to be monitored by a network monitoring system comprising a plurality of monitoring devices (stations 12) and a central measurement station 13. Each of the monitoring devices is associated with a respective one of the logical segments of the network. As will become clear below, each monitoring device is operative to randomly sample the packets on its associated network segment and transmit data on the sampled packets back to the measurement station 13 for processing and analysis.

The form of each monitoring device is illustrated in Figure 3. The device comprises a network interface 20, a microprocessor 21, and ROM (non-volatile, pre-programmed memory) and RAM (rewritable memory) units 22 and 23. These units 20 to 23 are all interconnected via address, data and control buses 27, 28 and 29. The network interface 20 is operative to carry out all the low level functions necessary to interface the monitoring device of Figure 3 to the network cable 10 and to pass received packets to a receive queue, in the form of a FIFO (First In First Out) buffer 25 in RAM 23. The

network interface is further operative to transmit packets held in a transmit queue, formed by a FIFO buffer 26, in RAM 23. The network interface 20 thus constitutes packet receive means and packet transmit means for the monitoring device. In the present example, the network interface 20 is arranged to receive all packets regardless of their destination address contained in the packet header. Furthermore, the network interface 20 is operative to pass only the header portion 30 of each received packet to the receive FIFO buffer 25.

The network interface 20 is arranged to operate in coordination with the microprocessor controller 21 and, in particular, informs the microprocessor 21 each time a packet header is inserted into the receive FIFO buffer 25, by means of a suitable interrupt control signal.

The network interface 20 also contains various counters 24 which hold a number of counts including the total number of packets received, the number of packets received which according to their CRC field are in error, the number of packets received below the minimum accepted length (RUNT packets), and the number of packets received above the maximum accepted length (JABBER).

Implementations of the network interface 20 for particular network protocols are well known in the art. Thus, for example, for an Ethernet network, the network interface 20 may be constituted by Intel Corporation chips 82502, 82501, and 82586; in this case an appropriate microprocessor constituting the microprocessor number 21 is the Intel processor 80186.

The ROM 22 holds the programs run by the microprocessor 21 and also a table of random count values predetermined according to an exponential distribution.

The processor 21 is operative to run a background program in which it does nothing (ie an idling program). The main working program for the processor 21 is an interrupt service routine which is called each time the network interface 20 generates a processor interrupt to tell the processor that it has stored a new packet header in the receive FIFO 25. The interrupt service routine, which will be described in more detail below, operates to randomly select a received packet header and form it into a collected-data packet together with the current count values of the counters 24; the random selection of received packet headers is effected by utilizing the predetermined random counts stored in ROM 22. The collected-data packet so formed is put into the transmit queue FIFO 26 and, in due course, is transmitted by the network interface 20 back to the measurement station 13. The header of each collected-data packet contains as its source address the address of the monitoring

device concerned while the destination address is that of the measurement station (alternatively, a multi-cast address can be used to which the measurement station is set to listen).

A more detailed description of the operation of the monitoring device will now be given with reference to Figure 4 which is a flow chart of the interrupt service routine run by the microprocessor 21. The microprocessor 21 will be taken to be in a state in which it is running its background (idling) program and in which it has one of the random count values held in an internal register (the fetching of the first count value upon switch-on of the monitoring device would be part of an initialization routine). It will also be assumed that the receive and transmit FIFO buffers 25 and 26 are empty.

On receiving a packet over the network cable 10, the network interface 20 passes the packet header to the receive FIFO buffer 25, updates its counters 24 and generates an interrupt signal for the microprocessor 21. On receipt of this interrupt, the microprocessor 21 executes the interrupt service routine illustrated in Figure 4. The first step 40 of this routine carries out the normal house-keeping tasks associated with such routines including saving the volatile environment parameters of the background program and masking further interrupts.

Next, the microprocessor decrements the random count value held in its internal register (step 41) and then checks the remaining value to see if this has been reduced to zero (step 42).

If the count value is still greater than zero, the microprocessor 21 discards the head entry in the receive FIFO buffer 25 (step 43).

Thereafter, the microprocessor must check the receive FIFO buffer 25 to see if any further packet headers have been entered into the buffer by the network interface 20 during the preceding steps of the interrupt service routine (step 44). Generally this will not be the case and the microprocessor will then exit its interrupt service routine and restore its background environment and unmask its interrupts (step 45). However, in the event that the receive FIFO buffer 25 contains a further packet header, the interrupt service routine will pass from step 44 back to step 41.

If during the test (step 42) carried out on the count value held in its internal register, the microprocessor 21 finds that this count value has been reduced to zero, the interrupt service routine will proceed to generate a collected-data packet 31 in respect of the packet header at the top of the receive FIFO buffer 25 (step 46). This collected-data packet 31 is assembled in the transmit FIFO buffer 26 from the received packet header 30, the count values from the counter 24, the address of the monitoring device (source address for the

collected-data packet) and the address of the measurement station (destination address for the collected-data packet header). After the collected-data packet has been assembled, the microprocessor 21 flags the network interface 20 to indicate that there is a packet ready for transmission. (The network interface 20 will transmit the packet as and when it is able and cancel the flag set by the microprocessor 21 once this has been done).

After completion of step 46 of the interrupt service routine, the microprocessor fetches a new random count from ROM 22 and stores this new random count in its internal register (step 47). The microprocessor then proceeds to step 44 and running of the interrupt service routine proceeds as previously described.

The size of the receive and transmit FIFO buffers 25 and 26 can be quite small, for example, sufficient to hold only two or three entries. This is possible with respect to the receive buffer 25 because in general the interval between packets received by the network interface 20 will be sufficient for the microprocessor 21 to run its interrupt service routine and clear the top entry from the receive buffer; in any event, the occasional overflowing of the receive buffer 25 is not of major consequence since the missing out of a packet will generally have minimal effect on the statistical measurements being conducted by the network monitoring system. This equally applies to the transmit buffer 26 where an overflow is even less likely to occur as its entries are only in respect of the randomly selected ones of the received packets.

The above-described implementation of the monitoring device does mean that the count values included in a collected-data packet from the counter 24 may not be the count values current at the time that the relevant packet was actually received by the network interface (this is because of the possible delay in actually processing the packet header). However, again, any discrepancy in this respect will be minor and will have minimal effect on the validity of the statistically determined results produced by the network monitoring system. Of course, it would be possible to design circuitry which associated the count values present in counters 24 with the header of each received packet; however, the added circuit complexity needed to do this is generally not justified.

The data structures used to implement the receive and transmit FIFO buffers 25 and 26 in RAM 23 will be apparent to a person skilled in the art and will therefore not be described herein. Furthermore, it will be appreciated that although in the Figure 3 embodiment the random selection of incoming packets has been effected by storing predetermined random numbers in ROM 22, these

random numbers could alternatively be generated as and when required by the processor 21 (although this is not preferred as it places extra processor requirements on the microprocessor).

Typically, the random numbers are such as to give an average skip between selected packets of ninety nine; other values may be more appropriate depending on traffic density, sampling period and acceptable statistical error level. The random selection of packets could be effected on a time basis rather than on the number of packets received.

The collected-data packets sent out by the monitoring devices 12 over the network are all received by the measurement station 13 which stores these packets and carries out subsequent processing and analysis.

There are a number of types of information that the measurement station 13 can derive from the packet samples provided by the collected-data packets:

a) Packet and error rates - Since packet and error counts are included in the collected-data packets, they can be tracked. Rates can easily be obtained by comparing counts in successive collected-data packets received from the same monitoring device 12.

b) Thresholds - Thresholds could be set on any of the values produced from the collected-data packets and higher level events generated. A typical threshold could be on the CRC rate, indicating that there may be a problem.

c) Traffic Matrices - By decoding the headers contained in the information fields of the collected-data packets many different traffic matrices can be obtained. The most basic traffic matrix would give the number of bytes and packets exchanged between each pair of stations on the network; such a matrix could be formed for each of the logical segments of the network as well as an overall matrix based on the maximum figures for each station pair found in the segment traffic matrices.

d) Address Mappings - The sampled packet headers generally contain information linking LAN addresses to higher level addresses. Thus where the LAN is an Ethernet LAN over which the TCP/IP protocol stack is being operated, the Ethernet to IP address mapping is easily obtained.

e) Availability - If the monitoring devices 12 are regarded as reliable, the lack of collected-data packets from a particular device or group of devices could be used to indicate network disconnections. In addition the lack of traffic from a particular station 11 (that of a file server, say) may indicate that it has failed.

It will of course be appreciated that a number of variations are possible to the described monitor-

ing device and network monitoring system. Thus for example, each collected-data packet formed by a monitoring device may contain data in respect of more than one randomly-selected packet. Furthermore, the data collected on a selected packet may comprise other elements to those described such as, for example, other fields of the selected packets as received by each monitoring device (in other words, fields additional to different from the packet header fields). The monitoring devices themselves can take the form of a stand-alone station as indicated in Figure 1, or cards slotted into existing network stations 11, or as part of the functionality provided by a bridge or router or as a process running on a processor of a connected network station.

Furthermore, the network monitoring device and system can be applied both to asynchronous datagram type networks such as Ethernet as well as to slotted networks where each station inserts data into a predetermined framed structure generated by a head station on the network.

Claims

1. A network monitoring device for monitoring the activity on a network carrying message packets each of which contains source and destination addresses, the monitoring device comprising receive means (20) for detecting and receiving message packets carried on the network, and processing means (21) operative to collect and process data about packets received by the receive means, this data including the source and destination addresses contained in the packet, characterized in that
 - the monitoring device (12) further comprises sampling means (21,22) for selecting some only of the packets detected by the receive means (20), and
 - the processing means (21) is operative to process data collected in respect of the selected packets only.
2. A network monitoring device according to claim 1, wherein the sampling means (21,22) effects its selection of packets based on the number of packets detected by the receive means (20).
3. A network monitoring device according to claim 1, wherein the sampling means (21) effects its selection of packets based on elapsed time since the previous packet selection.
4. A network monitoring device according to any one of claims 1 to 3, wherein the sampling means (21, 22) effects its selection of packets in a statistically random manner.
5. A network monitoring device according to claim 1, wherein the processing means (21) is operative to process the data collected on the selected packets to construct a traffic matrix.
6. A network monitoring device according to claim 1, wherein the monitoring device further comprises transmit means (20) operative to transmit message packets over a network, and wherein the processing means (21) is operative to process the data collected on said selected packets by forming collected-data packets (31) in each of which is included the data collected in respect of one or more of said selected packets, the processing means being operative to cause said transmit means (20) to transmit each said collected-data packet (31) for remote receipt and further processing.
7. A network monitoring device according to claim 6, further comprising counter means (24) for keeping a running count in respect of at least one of the following:
 - the total number of packets received by the receive means (20);
 - the total number of errored packets received by the receive means (20);
 - the total number of incorrectly-sized packets received by the receive means (20);
 the processing means (21) being operative to include the current value of the or each count in the data collected in respect of each selected packet whereby this count data is incorporated in the corresponding collected-data packet (31).
8. A network monitoring system operative to monitoring the activity on a network carrying message packets each of which contains source and destination addresses, the monitoring system comprising:
 - at least one network monitoring device (12) according to claim 6, with its receive means (20) connected to the network, and
 - a measurement station (13) connected to receive said collected-data packets (31) from the transmit means (20) of the or each said network monitoring device (12) and to extract and further process the data on the said selected packets that is contained in the collected-data packets (31).
9. A network monitoring system according to

claim 8, wherein the measurement station (13) is operative to construct traffic matrices from the data contained in the collected-data packets (31) received by the station.

5

10. A network monitoring system according to claim 8, wherein the network to be monitored includes one or more bridges and/or routers (14) separating the network into a plurality of logical segments and serving to restrict unnecessary packet flow between said segments, the monitoring system including a plurality of said network monitoring devices (12) each with its receive means (20) connected to a respective said segment of the network.

10

15

11. A network monitoring system according to any one of claims 8 to 10, wherein the measurement station (13) is connected to the same network (10) as the or each monitoring device (12), the transmit means (20) of each monitoring device (12) being operative to transmit said collected-data packets on said same network.

20

25

30

35

40

45

50

55

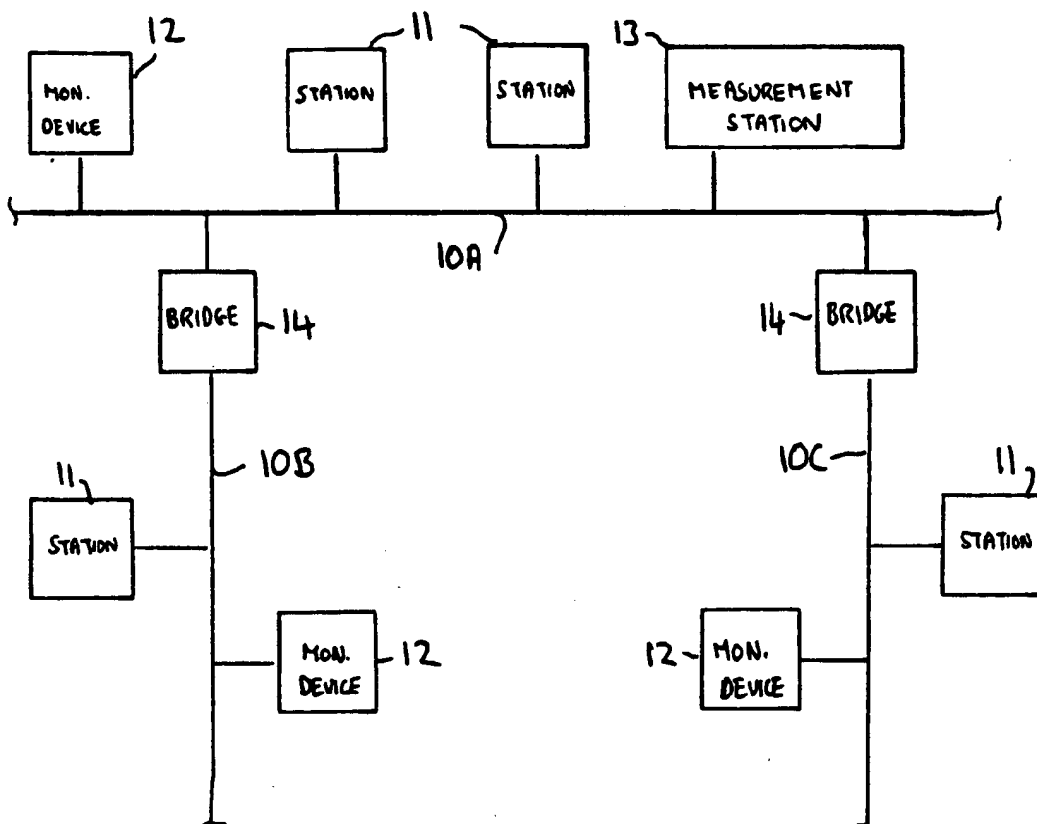


FIG. 1




FIG. 2




FIG. 3

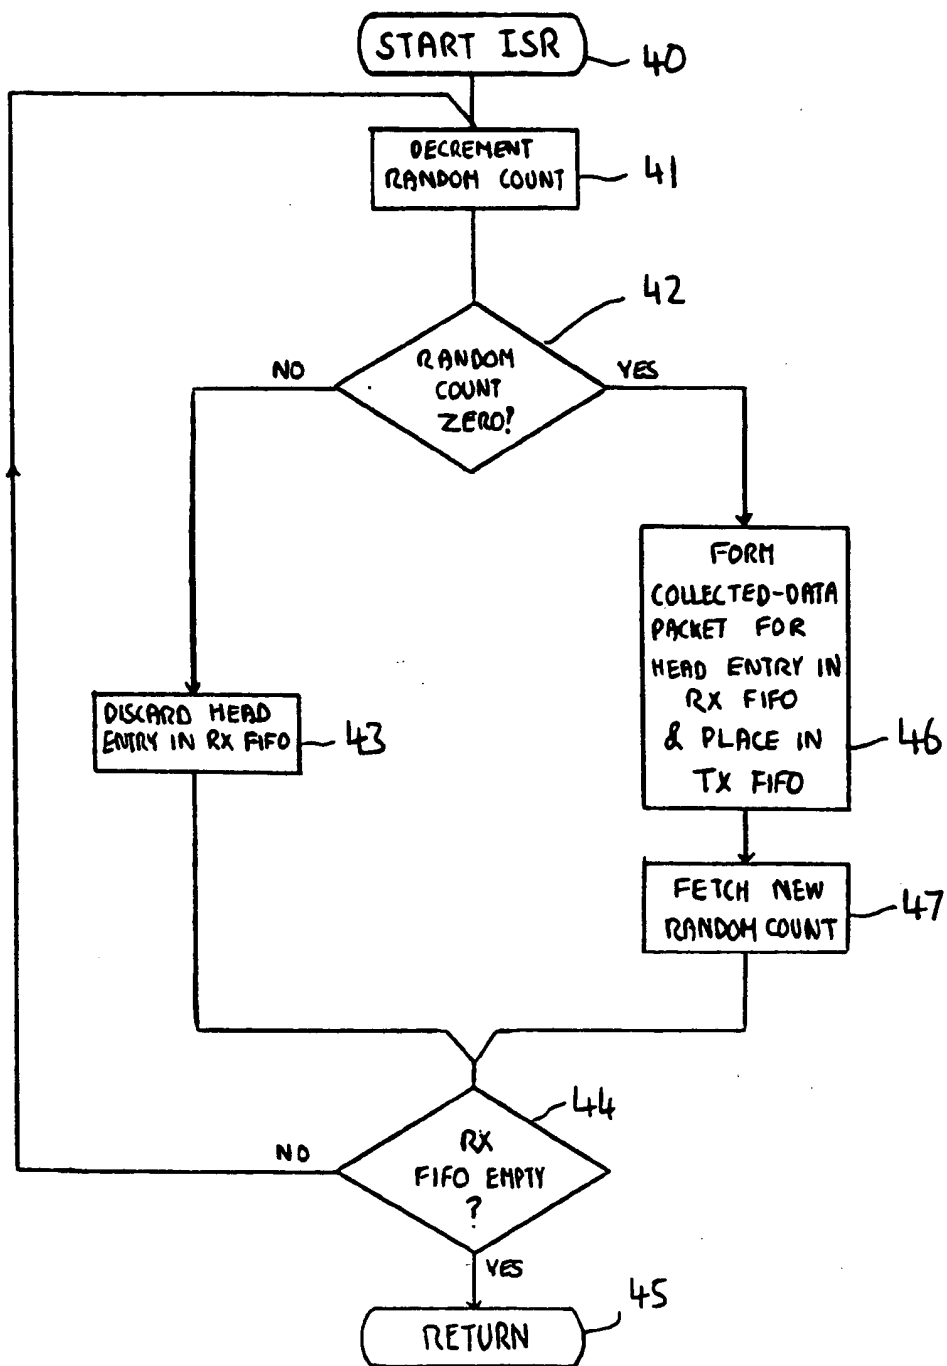


FIG. 4



European
Patent Office

EUROPEAN SEARCH REPORT

Application Number

EP 90 31 0699

DOCUMENTS CONSIDERED TO BE RELEVANT			
Category	Citation of document with indication, where appropriate, of relevant passages	Relevant to claim	CLASSIFICATION OF THE APPLICATION (Int. Cl.5)
A	IEEE NETWORK, vol. 1, no. 3, July 1987, pages 32-39; D. RITTER et al.: "A multi-purpose, distributed LAN traffic monitoring tool" * Page 32, right-hand column, lines 4-7; page 33, right-hand column, lines 19-30; page 34, left-hand column, lines 6-37; page 34, right-hand column, lines 7-33; page 35, right-hand column, lines 25-37 * - - -	1-11	H 04 L 12/26
A	IEEE NETWORK, vol. 1, no. 3, July 1987, pages 13-19; M. SOHA: "A distributed approach to LAN monitoring using intelligent high performance monitors" * Page 14, right-hand column, lines 6-21; page 15, left-hand column, lines 27-34; page 15, right-hand column, lines 10-38 * - - -	1,4-11	
A	SOFTWARE PRACTICE & EXPERIENCE, vol. 16, no. 7, July 1986, pages 671-687; S. VASSILADES et al.: "A monitor tool for a network based on the cambridge ring" * Page 675, lines 23-38; page 677, lines 33-37; page 680, lines 18-30 * - - - - -	1-3,7	
			TECHNICAL FIELDS SEARCHED (Int. Cl.5)
			H 04 L G 06 F
The present search report has been drawn up for all claims			
Place of search		Date of completion of search	Examiner
The Hague		19 June 91	DE LA FUENTE DEL AGU
CATEGORY OF CITED DOCUMENTS X: particularly relevant if taken alone Y: particularly relevant if combined with another document of the same category A: technological background O: non-written disclosure P: intermediate document T: theory or principle underlying the invention E: earlier patent document, but published on, or after the filing date D: document cited in the application L: document cited for other reasons ----- &: member of the same patent family, corresponding document			

THIS PAGE BLANK (USPTO)